

02-15-CU

A

Practitioner's Docket No. FORE-54**PATENT**JC690 U.S. PTO
02/14/00

Preliminary Classification:

Proposed Class:

Subclass:

NOTE: "All applicants are requested to include a preliminary classification on newly filed patent applications. The preliminary classification, preferably class and subclass designations, should be identified in the upper right-hand corner of the letter of transmittal accompanying the application papers, for example 'Proposed Class 2, subclass 129.'" M.P.E.P. § 601, 7th ed.

JC611 U.S. PTO
09/503673
02/14/00**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

Box Patent Application
Assistant Commissioner for Patents
Washington, D.C. 20231

NEW APPLICATION TRANSMITTAL

Transmitted herewith for filing is the patent application of

Inventor(s): Meenarachagan Vishnu

WARNING: 37 C.F.R. § 1.41(a)(1) points out:

"(a) A patent is applied for in the name or names of the actual inventor or inventors.

"(1) The inventorship of a nonprovisional application is that inventorship set forth in the oath or declaration as prescribed by § 1.63, except as provided for in § 1.53(d)(4) and § 1.63(d). If an oath or declaration as prescribed by § 1.63 is not filed during the pendency of a nonprovisional application, the inventorship is that inventorship set forth in the application papers filed pursuant to § 1.53(b), unless a petition under this paragraph accompanied by the fee set forth in § 1.17(i) is filed supplying or changing the name or names of the inventor or inventors."

For (title): METHOD AND APPARATUS FOR DYNAMIC BITMAP GENERATOR SCHEDULER

CERTIFICATION UNDER 37 C.F.R. § 1.10*

(Express Mail label number is mandatory.)

(Express Mail certification is optional.)

I hereby certify that this New Application Transmittal and the documents referred to as attached therein are being deposited with the United States Postal Service on this date February 14, 2000, in an envelope as "Express Mail Post Office to Addressee," mailing Label Number EL396485522US, addressed to the: Assistant Commissioner for Patents, Washington, D.C. 20231.

Tracey L. Milka

(type or print name of person mailing paper)

Tracey L. Milka

Signature of person mailing paper

WARNING: Certificate of mailing (first class) or facsimile transmission procedures of 37 C.F.R. § 1.8 cannot be used to obtain a date of mailing or transmission for this correspondence.

***WARNING:** Each paper or fee filed by "Express Mail" **must** have the number of the "Express Mail" mailing label placed thereon prior to mailing. 37 C.F.R. § 1.10(b).

"Since the filing of correspondence under § 1.10 without the Express Mail mailing label thereon is an oversight that can be avoided by the exercise of reasonable care, requests for waiver of this requirement will **not** be granted on petition." Notice of Oct. 24, 1996, 60 Fed. Reg. 56,439, at 56,442.

(New Application Transmittal [4-1]—page 1 of 11)

1. Type of Application

This new application is for a(n)

(check one applicable item below)

- ☒ Original (nonprovisional)
☐ Design
☐ Plant

WARNING: Do not use this transmittal for a completion in the U.S. of an International Application under 35 U.S.C. § 371(c)(4), unless the International Application is being filed as a divisional, continuation or continuation-in-part application.

WARNING: Do not use this transmittal for the filing of a provisional application.

NOTE: If one of the following 3 items apply, then complete and attach ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF A PRIOR U.S. APPLICATION CLAIMED and a NOTIFICATION IN PARENT APPLICATION OF THE FILING OF THIS CONTINUATION APPLICATION.

- ☐ Divisional.
☐ Continuation.
☐ Continuation-in-part (C-I-P).

2. Benefit of Prior U.S. Application(s) (35 U.S.C. §§ 119(e), 120, or 121)

NOTE: A nonprovisional application may claim an invention disclosed in one or more prior filed copending nonprovisional applications or copending international applications designating the United States of America. In order for a nonprovisional application to claim the benefit of a prior filed copending nonprovisional application or copending international application designating the United States of America, each prior application must name as an inventor at least one inventor named in the later filed nonprovisional application and disclose the named inventor's invention claimed in at least one claim of the later filed nonprovisional application in the manner provided by the first paragraph of 35 U.S.C. § 112. Each prior application must also be:

- (i) An international application entitled to a filing date in accordance with PCT Article 11 and designating the United States of America; or
- (ii) Complete as set forth in § 1.51(b); or
- (iii) Entitled to a filing date as set forth in § 1.53(b) or § 1.53(d) and include the basic filing fee set forth in § 1.16; or
- (iv) Entitled to a filing date as set forth in § 1.53(b) and have paid therein the processing and retention fee set forth in § 1.21(f) within the time period set forth in § 1.53(f).

37 C.F.R. § 1.78(a)(1).

NOTE: If the new application being transmitted is a divisional, continuation or a continuation-in-part of a parent case, or where the parent case is an International Application which designated the U.S., or benefit of a prior provisional application is claimed, then check the following item and complete and attach ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.

WARNING: If an application claims the benefit of the filing date of an earlier filed application under 35 U.S.C. §§ 120, 121 or 365(c), the 20-year term of that application will be based upon the filing date of the earliest U.S. application that the application makes reference to under 35 U.S.C. §§ 120, 121 or 365(c). (35 U.S.C. § 154(a)(2) does not take into account, for the determination of the patent term, any application on which priority is claimed under 35 U.S.C. §§ 119, 365(a) or 365(b).) For a c-i-p application, applicant should review whether any claim in the patent that will issue is supported by an earlier application and, if not, the applicant should consider canceling the reference to the earlier filed application. The term of a patent is not based on a claim-by-claim approach. See Notice of April 14, 1995, 60 Fed. Reg. 20,195, at 20,205.

(New Application Transmittal [4-1]—page 2 of 11)

WARNING: When the last day of pendency of a provisional application falls on a Saturday, Sunday, or Federal holiday within the District of Columbia, any nonprovisional application claiming benefit of the provisional application **must** be filed prior to the Saturday, Sunday, or Federal holiday within the District of Columbia. See 37 C.F.R. § 1.78(a)(3).

- ☐ The new application being transmitted claims the benefit of prior U.S. application(s). Enclosed are ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.

3. Papers Enclosed

A. Required for filing date under 37 C.F.R. § 1.53(b) (Regular) or 37 C.F.R. § 1.153 (Design) Application

22 Pages of specification

12 Pages of claims

10 Sheets of drawing

WARNING: **DO NOT** submit original drawings. A high quality copy of the drawings should be supplied when filing a patent application. The drawings that are submitted to the Office must be on strong, white, smooth, and non-shiny paper and meet the standards according to § 1.84. If corrections to the drawings are necessary, they should be made to the original drawing and a high-quality copy of the corrected original drawing then submitted to the Office. Only one copy is required or desired. For comments on proposed then-new 37 C.F.R. § 1.84, see Notice of March 9, 1988 (1990 O.G. 57-62).

NOTE: "Identifying indicia, if provided, should include the application number or the title of the invention, inventor's name, docket number (if any), and the name and telephone number of a person to call if the Office is unable to match the drawings to the proper application. This information should be placed on the back of each sheet of drawing a minimum distance of 1.5 cm. (5/8 inch) down from the top of the page . . ." 37 C.F.R. § 1.84(c)).

(complete the following, if applicable)

- ☐ The enclosed drawing(s) are photograph(s), and there is also attached a "PETITION TO ACCEPT PHOTOGRAPH(S) AS DRAWING(S)." 37 C.F.R. § 1.84(b).
- ☐ formal
- ☒ informal

B. Other Papers Enclosed

2 Pages of declaration and power of attorney

1 Pages of abstract

0 Other

4. Additional papers enclosed

- ☐ Amendment to claims
- ☐ Cancel in this applications claims _____ before calculating the filing fee. (At least one original independent claim must be retained for filing purposes.)
- ☐ Add the claims shown on the attached amendment. (Claims added have been numbered consecutively following the highest numbered original claims.)
- ☐ Preliminary Amendment
- ☐ Information Disclosure Statement (37 C.F.R. § 1.98)
- ☐ Form PTO-1449 (PTO/SB/08A and 08B)
- ☐ Citations

- ☐ Declaration of Biological Deposit
- ☐ Submission of "Sequence Listing," computer readable copy and/or amendment pertaining thereto for biotechnology invention containing nucleotide and/or amino acid sequence.
- ☐ Authorization of Attorney(s) to Accept and Follow Instructions from Representative
- ☐ Special Comments
- ☐ Other

5. Declaration or oath (including power of attorney)

NOTE: A newly executed declaration is not required in a continuation or divisional application provided that the prior nonprovisional application contained a declaration as required, the application being filed is by all or fewer than all the inventors named in the prior application, there is no new matter in the application being filed, and a copy of the executed declaration filed in the prior application (showing the signature or an indication thereon that it was signed) is submitted. The copy must be accompanied by a statement requesting deletion of the names of person(s) who are not inventors of the application being filed. If the declaration in the prior application was filed under § 1.47, then a copy of that declaration must be filed accompanied by a copy of the decision granting § 1.47 status or, if a nonsigning person under § 1.47 has subsequently joined in a prior application, then a copy of the subsequently executed declaration must be filed. See 37 C.F.R. §§ 1.63(d)(1)-(3).

NOTE: A declaration filed to complete an application must be executed, identify the specification to which it is directed, identify each inventor by full name including family name and at least one given name, without abbreviation together with any other given name or initial, and the residence, post office address and country or citizenship of each inventor, and state whether the inventor is a sole or joint inventor. 37 C.F.R. § 1.63(a)(1)-(4).

NOTE: "The inventorship of a nonprovisional application is that inventorship set forth in the oath or declaration as prescribed by § 1.62, except as provided for in § 1.53(d)(4) and § 1.63(d). If an oath or declaration as prescribed by § 1.63 is not filed during the pendency of a nonprovisional application, the inventorship is that inventorship set forth in the application papers filed pursuant to § 1.53(b), unless a petition under this paragraph accompanied by the fee set forth in § 1.17(i) is filed supplying or changing the name or names of the inventor or inventors." 37 C.F.R. § 1.41(a)(1).

☒ Enclosed

Executed by

(check all applicable boxes)

☒ inventor(s).

☐ legal representative of inventor(s).
37 C.F.R. §§ 1.42 or 1.43.

☐ joint inventor or person showing a proprietary interest on behalf of inventor who refused to sign or cannot be reached.

☐ This is the petition required by 37 C.F.R. § 1.47 and the statement required by 37 C.F.R. § 1.47 is also attached. See item 13 below for fee.

☐ Not Enclosed.

NOTE: Where the filing is a completion in the U.S. of an International Application or where the completion of the U.S. application contains subject matter in addition to the International Application, the application may be treated as a continuation or continuation-in-part, as the case may be, utilizing ADDED PAGE FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION CLAIMED.

☐ Application is made by a person authorized under 37 C.F.R. § 1.41(c) on behalf of all the above named inventor(s).

(The declaration or oath, along with the surcharge required by 37 C.F.R. § 1.16(e) can be filed subsequently).

- ☐ Showing that the filing is authorized.
(not required unless called into question. 37 C.F.R. § 1.41(d))

6. Inventorship Statement

WARNING: If the named inventors are each not the inventors of all the claims an explanation, including the ownership of the various claims at the time the last claimed invention was made, should be submitted.

The inventorship for all the claims in this application are:

- ☒ The same.

or

- ☐ Not the same. An explanation, including the ownership of the various claims at the time the last claimed invention was made,
☐ is submitted.
☐ will be submitted.

7. Language

NOTE: An application including a signed oath or declaration may be filed in a language other than English. An English translation of the non-English language application and the processing fee of \$130.00 required by 37 C.F.R. § 1.17(k) is required to be filed with the application, or within such time as may be set by the Office. 37 C.F.R. § 1.52(d).

- ☒ English
☐ Non-English
☐ The attached translation includes a statement that the translation is accurate. 37 C.F.R. § 1.52(d).

8. Assignment

- ☒ An assignment of the invention to FORE Systems, Inc.

- ☒ is attached. A separate ☒ "COVER SHEET FOR ASSIGNMENT (DOCUMENT) ACCOMPANYING NEW PATENT APPLICATION" or ☐ FORM PTO 1595 is also attached.
☐ will follow.

NOTE: "If an assignment is submitted with a new application, send two separate letters—one for the application and one for the assignment." Notice of May 4, 1990 (1114 O.G. 77-78).

WARNING: A newly executed "CERTIFICATE UNDER 37 C.F.R. § 3.73(b)" must be filed when a continuation-in-part application is filed by an assignee. Notice of April 30, 1993, 1150 O.G. 62-64.

(New Application Transmittal [4-1]—page 5 of 11)

001172025050

9. Certified Copy

Certified copy(ies) of application(s)

| Country | Appln. No. | Filed |
|---------|------------|-------|
| Country | Appln. No. | Filed |
| Country | Appln. No. | Filed |

from which priority is claimed

- ☐ is (are) attached.
☐ will follow.

NOTE: The foreign application forming the basis for the claim for priority must be referred to in the oath or declaration. 37 C.F.R. § 1.55(a) and 1.63.

NOTE: This item is for any foreign priority for which the application being filed directly relates. If any parent U.S. application or International Application from which this application claims benefit under 35 U.S.C. § 120 is itself entitled to priority from a prior foreign application, then complete item 18 on the ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED.

10. Fee Calculation (37 C.F.R. § 1.16)

A. ☒ Regular application

| CLAIMS AS FILED | | | |
|------------------------------|--------------|------------|----------------------------------|
| Number filed | Number Extra | Rate | Basic Fee 37 C.F.R. § 1.16(a) |
| | | | \$780.00 690.00 |
| Total | | | |
| Claims (37 C.F.R. § 1.16(c)) | 38 - 20 = 18 | × \$ 18.00 | 324.00 |
| Independent | | | |
| Claims (37 C.F.R. § 1.16(b)) | 10 - 3 = 7 | × \$ 78.00 | 546.00 |
| Multiple dependent claim(s), | | | |
| if any (37 C.F.R. § 1.16(d)) | | + \$260.00 | |

- ☐ Amendment cancelling extra claims is enclosed.
☐ Amendment deleting multiple-dependencies is enclosed.
☐ Fee for extra claims is not being paid at this time.

NOTE: If the fees for extra claims are not paid on filing they must be paid or the claims cancelled by amendment, prior to the expiration of the time period set for response by the Patent and Trademark Office in any notice of fee deficiency. 37 C.F.R. § 1.16(d).

Filing Fee Calculation \$ 1,560.00

B. ☐ Design application
(\$310.00—37 C.F.R. § 1.16(f))

Filing Fee Calculation \$

(New Application Transmittal [4-1]—page 6 of 11)

- C. ☐ Plant application
(\$480.00—37 C.F.R. § 1.16(g))

Filing fee calculation

\$ _____

11. Small Entity Statement(s)

- ☐ Statement(s) that this is a filing by a small entity under 37 C.F.R. § 1.9 and 1.27 is (are) attached.

WARNING: "Status as a small entity must be specifically established in each application or patent in which the status is available and desired. Status as a small entity in one application or patent does not affect any other application or patent, including applications or patents which are directly or indirectly dependent upon the application or patent in which the status has been established. The refiling of an application under § 1.53 as a continuation, division, or continuation-in-part (including a continued prosecution application under § 1.53(d)), or the filing of a reissue application requires a new determination as to continued entitlement to small entity status for the continuing or reissue application. A nonprovisional application claiming benefit under 35 U.S.C. § 119(e), 120, 121, or 365(c) of a prior application, or a reissue application may rely on a statement filed in the prior application or in the patent if the nonprovisional application or the reissue application includes a reference to the statement in the prior application or in the patent or includes a copy of the statement in the prior application or in the patent and status as a small entity is still proper and desired. The payment of the small entity basic statutory filing fee will be treated as such a reference for purposes of this section." 37 C.F.R. § 1.28(a)(2).

WARNING: "Small entity status must not be established when the person or persons signing the . . . statement can unequivocally make the required self-certification." M.P.E.P., § 509.03, 6th ed., rev. 2, July 1996 (emphasis added).

(complete the following, if applicable)

- ☐ Status as a small entity was claimed in prior application
_____/_____, filed on _____, from which benefit
is being claimed for this application under:

35 U.S.C. § ☐ 119(e),
☐ 120,
☐ 121,
☐ 365(c),

and which status as a small entity is still proper and desired.

- ☐ A copy of the statement in the prior application is included.

Filing Fee Calculation (50% of A, B or C above)

\$ _____

NOTE: Any excess of the full fee paid will be refunded if small entity status is established and a refund request are filed within 2 months of the date of timely payment of a full fee. The two-month period is not extendable under § 1.136. 37 C.F.R. § 1.28(a).

12. Request for International-Type Search (37 C.F.R. § 1.104(d))

(complete, if applicable)

- ☐ Please prepare an international-type search report for this application at the time when national examination on the merits takes place.

13. Fee Payment Being Made at This Time

☐ Not Enclosed

☐ No filing fee is to be paid at this time.

(This and the surcharge required by 37 C.F.R. § 1.16(e) can be paid subsequently.)

☒ Enclosed

☒ Filing fee \$ 1,560.00

☒ Recording assignment
(\$40.00; 37 C.F.R. § 1.21(h))
(See attached "COVER SHEET FOR
ASSIGNMENT ACCOMPANYING NEW
APPLICATION".) \$ 40.00

☐ Petition fee for filing by other than all the
inventors or person on behalf of the inventor
where inventor refused to sign or cannot be
reached
(\$130.00; 37 C.F.R. §§ 1.47 and 1.17(i)) \$ _____

☐ For processing an application with a
specification in
a non-English language
(\$130.00; 37 C.F.R. §§ 1.52(d) and 1.17(k)) \$ _____

☐ Processing and retention fee
(\$130.00; 37 C.F.R. §§ 1.53(d) and 1.21(l)) \$ _____

☐ Fee for international-type search report
(\$40.00; 37 C.F.R. § 1.21(e)) \$ _____

NOTE: 37 C.F.R. § 1.21(l) establishes a fee for processing and retaining any application that is abandoned for failing to complete the application pursuant to 37 C.F.R. § 1.53(f) and this, as well as the changes to 37 C.F.R. §§ 1.53 and 1.78(a)(1), indicate that in order to obtain the benefit of a prior U.S. application, either the basic filing fee must be paid, or the processing and retention fee of § 1.21(l) must be paid, within 1 year from notification under § 53(f).

Total fees enclosed \$ 1,600.00

14. Method of Payment of Fees

☒ Check in the amount of \$ 1,560.00 & 40.00

☐ Charge Account No. _____ in the amount of
\$ _____

A duplicate of this transmittal is attached.

NOTE: Fees should be itemized in such a manner that it is clear for which purpose the fees are paid. 37 C.F.R. § 1.22(b).

15. Authorization to Charge Additional Fees

WARNING: If no fees are to be paid on filing, the following items should not be completed.

WARNING: Accurately count claims, especially multiple dependent claims, to avoid unexpected high charges, if extra claim charges are authorized.

- ☒ The Commissioner is hereby authorized to charge the following additional fees by this paper and during the entire pendency of this application to Account No. 19-0737:

- ☒ 37 C.F.R. § 1.16(a), (f) or (g) (filing fees)
☒ 37 C.F.R. § 1.16(b), (c) and (d) (presentation of extra claims)

NOTE: Because additional fees for excess or multiple dependent claims not paid on filing or on later presentation must only be paid or these claims cancelled by amendment prior to the expiration of the time period set for response by the PTO in any notice of fee deficiency (37 C.F.R. § 1.16(d)), it might be best not to authorize the PTO to charge additional claim fees, except possibly when dealing with amendments after final action.

- ☐ 37 C.F.R. § 1.16(e) (surcharge for filing the basic filing fee and/or declaration on a date later than the filing date of the application)
☐ 37 C.F.R. § 1.17(a)(1)-(5) (extension fees pursuant to § 1.136(a)).
☐ 37 C.F.R. § 1.17 (application processing fees)

NOTE: ". . . A written request may be submitted in an application that is an authorization to treat any concurrent or future reply, requiring a petition for an extension of time under this paragraph for its timely submission, as incorporating a petition for extension of time for the appropriate length of time. An authorization to charge all required fees, fees under § 1.17, or all required extension of time fees will be treated as a constructive petition for an extension of time in any concurrent or future reply requiring a petition for an extension of time under this paragraph for its timely submission. Submission of the fee set forth in § 1.17(a) will also be treated as a constructive petition for an extension of time in any concurrent reply requiring a petition for an extension of time under this paragraph for its timely submission." 37 C.F.R. § 1.136(a)(3).

- ☐ 37 C.F.R. § 1.18 (issue fee at or before mailing of Notice of Allowance, pursuant to 37 C.F.R. § 1.311(b))

NOTE: Where an authorization to charge the issue fee to a deposit account has been filed before the mailing of a Notice of Allowance, the issue fee will be automatically charged to the deposit account at the time of mailing the notice of allowance. 37 C.F.R. § 1.311(b).

NOTE: 37 C.F.R. § 1.28(b) requires "Notification of any change in status resulting in loss of entitlement to small entity status must be filed in the application . . . prior to paying, or at the time of paying, . . . the issue fee. . . ." From the wording of 37 C.F.R. § 1.28(b), (a) notification of change of status must be made even if the fee is paid as "other than a small entity" and (b) no notification is required if the change is to another small entity.

16. Instructions as to Overpayment

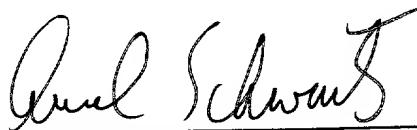
NOTE: "... Amounts of twenty-five dollars or less will not be returned unless specifically requested within a reasonable time, nor will the payer be notified of such amounts; amounts over twenty-five dollars may be returned by check or, if requested, by credit to a deposit account." 37 C.F.R. § 1.26(a).

- ☒ Credit Account No. 19-0737
- ☐ Refund

Reg. No. 30,587

Tel. No. (412) 621-9222

Customer No.



SIGNATURE OF PRACTITIONER

Ansel M. Schwartz

(type or print name of attorney)

One Sterling Plaza
201 N. Craig Street

P.O. Address

Suite 304
Pittsburgh, PA 15213

☒ **Incorporation by reference of added pages**

(check the following item if the application in this transmittal claims the benefit of prior U.S. application(s) (including an international application entering the U.S. stage as a continuation, divisional or C-I-P application) and complete and attach the ADDED PAGES FOR NEW APPLICATION TRANSMITTAL WHERE BENEFIT OF PRIOR U.S. APPLICATION(S) CLAIMED)

- ☐ Plus Added Pages for New Application Transmittal Where Benefit of Prior U.S. Application(s) Claimed

Number of pages added _____

- ☐ Plus Added Pages for Papers Referred to in Item 4 Above

Number of pages added _____

- ☐ Plus added pages deleting names of inventor(s) named in prior application(s) who is/are no longer inventor(s) of the subject matter claimed in this application.

Number of pages added _____

- ☒ Plus "Assignment Cover Letter Accompanying New Application"

Number of pages added _____ 4

☐ **Statement Where No Further Pages Added**

(if no further pages form a part of this Transmittal, then end this Transmittal with this page and check the following item)

- ☐ This transmittal ends with this page.

METHOD AND APPARATUS FOR DYNAMIC BITMAP GENERATOR SCHEDULER

FIELD OF THE INVENTION

The present invention is related to a bitmap scheduler.
5 More specifically, the present invention is related to a
hierarchical bitmap generator scheduler.

BACKGROUND OF THE INVENTION

Service scheduling is the primary mechanism for providing
Quality of Service (QoS) guarantees on a per-VC basis in
10 Asynchronous Transfer Mode (ATM) networks. Such service scheduling
schemes must satisfy a number of requirements in order to be used
in practical ATM switches and multiplexers. Firstly, such a
service scheduling scheme must guarantee specified service rate to
each virtual connection (VC), irrespective of the traffic patterns
15 in the VCS. Secondly, the scheduling scheme must flexibly allocate
excess (i.e., temporarily unused and unallocated) bandwidth among
the active VCs. Thirdly, the outgoing traffic streams of each VC
and VPs must be smooth (shaped) and not bursty. Fourthly, the
service rate given to a VC or a group of VCs must not exceed a
20 specified upper bound. Most importantly, the scheduling algorithm
must be simple so that the scheduling decision can be performed
using only a few operations per cell time.

Previously proposed schemes such as the Weighted Round
Robin (WRR), Packetized Generalized Processor Sharing (PGPS) [A. K.
25 Parekh and R. G. Gallager. A generalized processor sharing
approach to flow control in integrated services networks: The
single node case. *IEEE/ACM Transactions on Networking*,
1(3):344-357, June 1993; S. Demers, A. Keshav and S. Shenker.
Analysis and simulation of a fair queuing algorithm. *Internet*

Research and Experience, 1, 1990, incorporated by reference herein], Self-Clocked Fair Queueing (SCFQ) [S. J. Golestani. A self-clocked fair queueing scheme for broadband applications. In *Proceedings of IEEE INFOCOM*, pages 636-646, June 1994, incorporated
5 by reference herein], Worst Case Fair Weighted Fair Queueing (WF²Q) [C. R. Bennet and H. Zhang. WF²Q: Worst-case fair weighted fair queueing. In *Proceedings of IEEE INFOCOM*, pages 120-128, 1996, incorporated by reference herein], and Virtual Clock [L. Zhang. Virtualclock: A new traffic control algorithm for packet switched
10 networks. *ACM Transactions on Computer Systems*, 9(2):101-124, May 1991, incorporated by reference herein] have either fallen short of these goals or are too complex to be implemented in high speed hardware cost-effectively.

SUMMARY OF THE INVENTION

15 The present invention pertains to a scheduler for a server. The scheduler comprises a first level generator associated with groups of connections. The scheduler comprises a second level generator associated with connections corresponding to the groups of connections. The first level generator identifying which
20 connections in the second level generator corresponds to a group in the first level generator that is to be considered for service. The second level generator identifies the connections corresponding to the group to receive service from the server. The second level generator in connection with the first level generator.

25 The present invention pertains to a method for scheduling service of a server. The method comprises the steps of identifying a group of connections with a first level generator to receive service from the server. Then there is the step of identifying

connections corresponding with the group of connections with a second level generator to receive service from the server.

The present invention pertains to an apparatus for serving connections. The apparatus comprises a server. The apparatus comprises a memory in which data of the connections is stored. The memory is connected to the server. The apparatus comprises a hierarchical scheduler connected to the server which schedules when the data of the connections in the memory is to receive service from the server. The scheduler is connected to the server and the memory.

The present invention pertains to an apparatus for serving connections. The apparatus comprises a server. The apparatus comprises a memory in which cells of the connections are stored. The memory is connected to the server. The apparatus comprises a scheduler connected to the server which schedules when the cells of the connections in the memory are to receive service from the server based on intercell intervals, wherein an intercell interval is how long the server takes to service a cell. The scheduler is connected to the server and the memory.

BRIEF DESCRIPTION OF THE DRAWINGS

In the accompanying drawings, the preferred embodiment of the invention and preferred methods of practicing the invention are illustrated in which:

Figure 1 shows weighted Round Robin weights.

Figure 2 shows a bitmap array.

Figure 3 shows dynamic bitmap generation.

Figure 4 shows a dynamic bitmap hierarchy.

Figure 5 is a schematic representation of a hierarchical dynamic bitmap generator scheduler.

Figure 6 is a schematic representation of a level 1
5 bitmap generator.

Figure 7 is a schematic representation of a level 1 filter-encoder.

Figure 8 is a schematic representation of a counter.

Figure 9 is a graph of how to schedule overbookable and
10 guaranteed bandwidth.

Figure 10 is a schematic representation of a scheduler of the present invention.

Figure 11 is a schematic representation of an apparatus of the present invention.

15 Figure 12 is a schematic representation of an apparatus of the present invention.

DETAILED DESCRIPTION

Referring now to the drawings wherein like reference numerals refer to similar or identical parts throughout the several
20 views, and more specifically to figure 10 thereof, there is shown a scheduler 10 for a server 12. The scheduler 10 comprises a first level generator 14 associated with groups of connections. The scheduler 10 comprises a second level generator 16 associated with

connections corresponding to the groups of connections. The first level generator 14 identifying which connections in the second level generator 16 corresponds to a group in the first level generator 14 that is to be considered for service. The second level generator 16 identifies the connections corresponding to the group to receive service from the server 12. The second level generator 16 in connection with the first level generator 14.

Preferably, the scheduler 10 includes a first level filter mechanism 18 which filters out inactive groups of connections. The first level filter mechanism 18 is connected to the first level generator 14 and the second level generator 16. The scheduler 10 preferably includes a second level filter mechanism 20 which filters out inactive connections. The second level filter mechanism 20 is connected to the second level generator 16. Preferably, the scheduler 10 includes a zero level generator 22 associated with supergroups corresponding with groups. The zero level generator 22 in connection with the first level generator 14. The zero level generator 22 identifying which groups in the first level generator 14 correspond to a supergroup in the zero level generator 22 that are considered for service.

The scheduler 10 preferably includes a zero level filter mechanism 23 which filters out inactive supergroups. The zero level filter mechanism 23 is connected to the zero level generator 22 and the first level generator 14. Preferably, the zero level generator 22 includes a zero level bitmap generator 24 which generates a zero level schedule bitmap which indicates the supergroup to be scheduled for service, the first level generator 14 includes a first level bitmap generator 26 which indicates the group to be scheduled for service, and the second level generator 16 includes a second level bitmap generator 28 which generates a

second level schedule bitmap which indicates the connections to be scheduled for service.

The zero level, first level and second level filter mechanism 20 preferably includes a zero level filter encoder 30, first level filter encoder 32 and second level filter encoder 34, respectively, which filters out inactive supergroups from the zero level schedule bitmap and encodes the zero level schedule bitmap with inactive supergroups removed, which filters out inactive groups from the first level schedule bitmap and encodes the first level schedule bitmap with inactive groups removed, and which filters out inactive connections from the second level schedule bitmap and encodes the second level schedule bitmap with inactive connections removed, respectively. Preferably, the scheduler 10 includes an interface 36 which maintains a zero level active bitmap 38, a first level active bitmap 40 and a second level active bitmap 42 having only active connections corresponding to the zero level schedule bitmap, first level schedule bitmap and second level schedule bitmap, respectively. Preferably, each active bitmap has a bit which is set to 1 when an associated connection is active and is set to 0 when an associated connection is inactive.

The zero level filter encoder 30 preferably reads the zero level schedule bitmap and ANDS it with the zero level active bitmap 38 to filter out inactive supergroups, the first level filter encoder 32 reads the first level schedule bitmap and ANDS it with the first level active bitmap 40 to filter out inactive groups, and the second level filter encoder 34 reads the second level schedule bitmap and ANDS it with the second level active bitmap 42 to filter out inactive supergroups. Preferably, the zero level bitmap generator 24, first level bitmap generator 26 and second level bitmap generator 28 dynamically generates bits for each supergroup, group and connection, respectively.

The zero level bitmap generator 24 preferably includes a counter 44 for each supergroup which is decremented as a function of an intercell interval, wherein the intercell interval is the time it takes for the server 12 to service a cell, the first level
5 bitmap generator 26 includes a counter 44 for each group which is decremented as a function of the intercell interval, and the second level bitmap generator 28 includes a counter 44 for each connection which is decremented as a function of the intercell interval. Preferably, the zero level bitmap generator 24 sets a bit for a
10 supergroup whose counter 44 decrements to zero, the first level bitmap generator 26 sets a bit for a group whose counter 44 decrements to zero, and the second level bitmap generator 28 sets a bit for a connection whose counter 44 decrements to zero. Alternatively, each counter at each level has a different number of
15 bits.

The zero level bitmap generator 24, first level bitmap generator 26 and second level bitmap generator 28 each preferably include a rate limiting counter 46 associated with each counter 44, wherein the bit for the supergroup, group or connection,
20 respectively, is set whenever both the counter 44 and the corresponding rate limiting counter 46 decrements to zero. Preferably, the zero level bitmap generator 24, first level bitmap generator 26 and second level bitmap generator 28 each generate a guaranteed rate bitmap for supergroups, groups and connections,
25 respectively, which receive service before any other supergroups, groups or connections, respectively, in the respective schedule bitmaps. The zero level bitmap generator 24, first level bitmap generator 26 and second level bitmap generator 28 preferably proportionately reduce the service to each supergroup, group and
30 connection, respectively, when overbooking occurs.

Preferably, connections arise from entities, and alternatively, the apparatus includes multiple counters associated with each entity which have multiple bits, including multiple schedule bitmaps associated with each entity that are used to
5 schedule connections from the corresponding entity at different priorities or a combination of priorities.

The present invention pertains to a method for scheduling service of a server 12. The method comprises the steps of identifying a group of connections with a first level generator 14
10 to receive service from the server 12. Then there is the step of identifying connections corresponding with the group of connections with a second level generator 16 to receive service from the server 12.

Preferably, after the identifying the group of connections step, there is the step of filtering out inactive groups of connections in regard to the first level generator 14. After the identifying the connections step, there is the step of filtering out inactive connections in regard to the second level generator 16. Preferably, before the step of identifying the group
15 of connections, there is the step of identifying groups in the first level generator 14 corresponding to a supergroup and a zero level generator 22.

After the identifying groups step, there is preferably the step of filtering out inactive supergroups of connections in regard to the zero level generator 22. Preferably, the filtering out the inactive supergroups step includes the step of ANDing a zero level schedule bitmap of the zero level bitmap generator 24 with a zero level active bitmap 38 of an interface 36 to filter out inactive supergroups. The filtering out the inactive groups step
25 preferably includes the step of ANDing a first level schedule
30

bitmap of the first level bitmap generator 26 with a first level active bitmap 40 of an interface 36 to filter out inactive groups.

Preferably, the filtering out the inactive connections step includes the step of ANDing a second level schedule bitmap of the second level bitmap generator 28 with a second level active bitmap 42 of an interface 36 to filter out inactive connections. The identifying the groups of connections step preferably includes the step of generating dynamically the zero level schedule bitmap, the identifying the group step includes the step of generating dynamically the first level schedule bitmap, and the identifying the connections step includes the step of generating dynamically the second level generator 16 schedule bitmap. Preferably, the step of generating the zero level schedule bitmap includes the step of decrementing a counter 44 for each supergroup every intercell interval; the step of generating the first level schedule bitmap includes the step of decrementing a counter 44 for each group every intercell interval; the step of generating the second level schedule bitmap includes the step of decrementing a counter 44 for each connection every intercell interval.

The present invention pertains to an apparatus 47 for serving connections, as shown in figure 11. The apparatus 47 comprises a server 12. The apparatus 47 comprises a memory 48 in which data of the connections is stored. The memory 48 is connected to the server 12. The apparatus 47 comprises a hierarchical scheduler 50 connected to the server 12 which schedules when the data of the connections in the memory 48 is to receive service from the server 12. The scheduler 50 is connected to the server 12 and the memory 48.

The present invention pertains to an apparatus 49 for serving connections, as shown in figure 12. The apparatus 49

comprises a server 12. The apparatus 49 comprises a memory 48 in which cells of the connections are stored. The memory 48 is connected to the server 12. The apparatus 49 comprises a scheduler 52 connected to the server 12 which schedules when the cells of the connections in the memory 48 are to receive service from the server 12 based on intercell intervals, wherein an intercell interval is how long the server 12 takes to service a cell. The scheduler 52 is connected to the server 12 and the memory 48.

Preferably, the intercell intervals are inversely proportional to bandwidth allocated to a connection. Spacing at intercell intervals of cells is performed preferably by either statically storing a set of schedule bitmaps or by dynamically generating the schedule bitmap specifying which connections are to be served.

In the operation of the preferred embodiment, the scheduler 10 describes a service scheduling scheme and its implementation for high-speed ATM switches and multiplexers. The scheduling scheme satisfies all of the following required properties. Moreover, the scheduling scheme can be implemented in high-speed hardware cost-effectively.

Minimum Specified Bandwidth Guarantee. Once a VC is admitted, an ATM scheduling scheme must guarantee a minimum specified bandwidth to each VC, irrespective of the traffic streams sharing the link. This is crucial for the ATM networks to guarantee specified QoS such as bounds on cell delay and cell loss rate on a per-VC basis.

Hierarchical Shaping. The outgoing VC streams must be smooth and not bursty. Bursty VC streams require larger buffer

space in downstream nodes and increase both the cell loss rate and the cell delay variation. Hierarchical shaping is desirable when VPs are considered as a single entity in downstream switching nodes.

5 **Hierarchical Rate Limiting.** In some scenarios, the service rate of a VC or a group of VCS must be upper bounded, as well. For example, if the VC or the group of VCS passes through a leased line of limited bandwidth, then the VC or VP needs to have an upper bound on the bandwidth it receives at the switch port.

10 **Overbooking.** Several service providers like to overbook their lines because they observe that their lines are usually underutilized. From a scheduler 10 point of view, overbooking means that the sum of the bandwidth of the admitted VCS can be greater than the link bandwidth and when there is congestion, the link bandwidth be shared proportional to the requested bandwidth.

15 **Overbooking with Minimum Guarantees.** Degradation in service rate with overbooking may not be acceptable to some time sensitive services such as CBR and rt-VBR. Such VCS must be guaranteed their specified bandwidth, while other VCS overbook
20 their bandwidth.

Flexible and Dynamic Adjustment of Excess Bandwidth Allocation. It is desirable to dynamically adjust the allocated bandwidth. This is useful, for example, to change the bandwidth allocation to Available Bit Rate (ABR) VCS depending on the
25 computed explicit rate (ER) values.

Fast VC Setup/Teardown. Initializing the scheduler 10 at VC setup/VC teardown must not involve more than few accesses of the memory 48 mapped registers of the scheduler 10. If the scheduler 10 requires the initialization of large data structures when VCS are setup/torn down, then the setup/teardown time is considerably increased.

To better understand the scheme, consider a one-dimensional array of WRR scheduler 10 weights indexed by the VC number as shown in figure 1. The software calculates the weight of the i th VC as follows:

$$w_i = \frac{MAX_WEIGHT \times r_i}{R} \quad (1)$$

where MAX_WEIGHT is 256 and r_i is the bandwidth requested by the i th VC and R is the line rate.

As noted above, one of the problems with the WRR scheduler 10 is that the outgoing VC streams are bursty, because WRR sends bursts of w_i cells from the i th VC. One way to make the outgoing VC streams smooth is to replace the one-dimensional array with a two-dimensional array of bits as shown in figure 2.

Suppose that the i th VC had a weight of w_i . Divide 256 by w_i to obtain its inter-cell interval, D_i in slots. Set every D_i th bit of the i th row to 1 and the remaining bits of the row to 0.

The operation of this bitmap scheduler 10 is as follows: The column 0 of the bits is first read and the VCS corresponding to

the bits set to '1' are served. This is followed by the column 1, 2, ... up to column 255. The cycle again starts with column 0.

There are at least three problems with this solution: The first is that it requires enormous amount of memory 48. To support 5 256K VCS, 256K x 256 bits of RAM is required. (Assuming a width of 256 bits is sufficient. Currently switch software sets the maximum weight in the WRR scheduler 10 to 256). Secondly, VC setup/teardown requires 256 memory 48 writes which will inevitably slow down VC setup/teardown times. The third problem is that the scheduler 10 may be spending time reading large numbers of empty VCS.

The first problem can be solved by dynamically generating bits for each VC rather than storing precomputed bit patterns. The bit pattern corresponding to the i th VC is simple: Each D_i th bit is set. This can be done by having a down counter as shown in figure 3. The counter 44 corresponding to the i th VC is loaded with the value D_i and the i th VC bit is set when the counter 44 counts down to zero, at which point the counter 44 is reloaded with the value of D_i .

20 This solution also solves the second problem, because now the i th VC setup only requires the initialization of the intercell interval, D_i . That is, it requires only one memory 48 access.

To solve the third problem and to avoid the large number of counters 44, registers and logic needed, the bitmaps are 25 organized as a hierarchy as shown in figure 4.

The Hierarchical Dynamic Bitmap Generator (HDBMG) is shown in figure 5. The scheduler 10 consists of the following

components: three Bitmap Generators (BMGs), three Filter-Encoder (FEs), Trident Interface (TI) and AD Bus Controller. The VCS are organized into a three level hierarchy, consisting of 256K (64 x 64 x 64) VCS, 4K (64 x 64) VC groups and 64 VC supergroups.

5 In other words, i th VC group consists of VCS $64i$ to $64i + 63$ and j th VC supergroup consists of VC groups $64j$ to $64j + 63$.

The level 0 (level 1, level 2) BMG generate schedule bitmaps which indicate the VC supergroup (VC groups, VCS) to be scheduled at each slot. The bitmap is generated using the intercell interval (D) and the current counter (C) values stored in registers (internal RAMs, external RAMs). The level 0 (level 1, level 2) FE filters out inactive VC supergroups and encodes the resulting bitmap to determine VC supergroups (VC groups, VCS) to be scheduled.

10 The bitmap generated by level 0 BMG is placed into a FIFO within the BMG. The level 0 FE pops the bitmaps, filters out inactive VC supergroups and encodes the resulting bitmap into a list of VC supergroup numbers as shown in figure 5. The level 1 BMG obtains the next VC supergroup number from the level 0 FE, 15 reads the corresponding data (intercell interval and current counter values) from the internal RAMs, generates the VC group bitmap and stores it in a FIFO. The level 1 FE pops the next bitmap, filters out inactive VC groups, and encodes the resulting bitmap into a list of VC groups. Similarly, the level 2 generates 20 VC bitmaps and level 2 FE filters out inactive VCS and encode the bitmaps into a list of VCS which are sent to the Trident Interface.

The BMGs generates schedule bitmaps which indicate which VCS (VC groups or VC supergroups) are scheduled at the current slot. The level 0 BMG is the simplest. It does not need any

external RAM because it only handles a single set of 64 VC supergroups. The data (D and C) can be stored in registers within the level 0 BMG. The operation of the level 0 BMG. The operation of the level 0 BMG is as follows: At each clock cycle, if its
5 bitmap FIFO is not full, it decrements all the 64 counters. If any of them have reached zero, then those counters are reloaded with the corresponding value of D. Also, the bits corresponding these VC supergroups are set to 1 in the schedule bitmap and pushed into the FIFO. Note that the bits corresponding to the VC supergroups
10 whose counters have not reached zero are set to 0.

The operation of the level 1 BMG is slightly more complex (see figure 6). As described above, the bitmaps generated by the level 0 BMG are placed in a FIFO. The level 0 filter-encoder (FE) pops these bitmaps, filters out inactive VC supergroups and encodes
15 the set bits of the resulting bitmap into a list of VC supergroups. The details of the operation of the FE are described below. The level 1 BMG has a similar organization as the level 0 BMG. In addition, it is connected to RAMs which contains the interval (D) and the counter values of the 4K (64 x 64) VC groups. The
20 operation of the level 1 BMG is as follows: A level 1 BMG requests and gets the next VC supergroup number from the level 0 FE. It reads the set of 64 D and C values of the 64 VC groups belong to the received VC supergroup. The level 1 BMG then computes the schedule bitmap and puts it in a FIFO. In fact, a (64 + 6)-bit
25 wide word containing the 64-bit wide bitmap plus the 6-bit wide VC supergroup number is put into the FIFO. Note that the operations such as accessing the next VC supergroup number, accessing data from the RAMs, computing the bitmap and writing back the updated counter values to the RAM can be pipelined to generate a schedule
30 bitmap at every clock cycle (as long as the FIFO is not full).

The operation of the level 2 BMG is almost identical to the level 1 BMG. The RAMs connected to the level 2 BMGS are considerably larger, containing the interval and counter values of the 256K (64 x 64 x 64) VCS. The level 1 FE pops the level 1 BMGS
5 FIFO, filters out inactive VC groups and encodes the resulting bitmap into a list of VC groups. The complete VC group number is obtained by concatenating the 6-bit VC supergroup number attached to the bitmap with the 6-bit indicating the position of the bit in the bitmap. That is, the output of the level 1 FE, is a list of
10 12-bit wide words indicating which of the 4K VC groups are to be scheduled. The level 2 BMG requests and received these 12-bit VC group number, reads the corresponding set of intervals and counter values of the 64 VCS, generates the bitmap and puts the bitmap into a FIFO. As done by the level 1 BMG, the level 2 BMG also attaches
15 the 12-bit VC group number with the generated bitmap. Therefore, the total width of the words pushed into the level 2 BMGs FIFO is (64 + 12 = 75)-bits.

The level 2 FE, pops these bitmaps, filters out inactive VCS and encode the bitmaps into a list of VCS to be scheduled.
20 Note that the 18-bit numbers indicating subset of possible 256K VCS are obtained by concatenating 12-bit VC group numbers attached to the bitmap with the 6-bit indicating the position of the bit in the bitmap.

The scheduler 10 employs three filter-encoders (FEs).
25 The level 0 FE pops the bitmap FIFO of level 0 BMG, filters out inactive VC supergroups from the bitmap and encodes the resulting bitmap into a list of VC supergroup numbers. As described in Subsection 3.5, the Trident Interface maintains a hierarchical bitmaps of the active VCS, i.e., VCS which have cells in their
30 per-VC queues. These bitmaps are referred to as active bitmaps to avoid confusion with the schedule bitmaps. The level 0 (level 1,

level 2) FE reads the corresponding active bitmaps and ANDs it with the schedule bitmaps to filter out inactive VC supergroups (VC groups, VCS).

Figure 7 shows the level 1 FE. It requests and gets the VC group schedule bitmap from the level 1 BMG. It then reads the active bitmap of the VC supergroup. Note that the VC supergroup number is attached to the bitmap. The FE filters out inactive VC groups by ANDing the schedule bitmap with the active bitmap. Finally, a priority encoder converts the most significant active bit of the resulting bitmap to VC group number. When the *get next VC group number* signal is asserted by the level 2 BMG, the most significant active bit is cleared and the priority encoder encodes the next most significant active bit. If there are no more active bits, then, a new set of bitmaps is loaded and used to encode the next VC group number.

The Trident Interface (TI) forms the interface 36 to the Trident ASIC. It receives a cell arrival information from the Trident ASIC and maintains the three level hierarchical active bitmaps. The Trident ASIC informs the port to schedule and the TI gets the next VC from the level 2 FE and sends it to the Trident ASIC. Trident products are available from FORE Systems, Inc., Warrendale, Pennsylvania.

The AD bus controller provides an interface 36 to the Switch Control Processor (SCP) to access internal registers, internal and external RAMS to set up and tear down VCS. The control and status registers are memory-mapped to the address space of the SCP.

As mentioned above, it is desirable to have two types of bandwidth allocations, overbookable bandwidth and guaranteed bandwidth. That is, the bandwidth of a VC is specified as the following 3-tuples (r_g, r_e, r_{\max}) , where r_g is the bandwidth to be
5 guaranteed, r_e is the additional excess bandwidth to be allocated (subject to availability) and r_{\max} is the maximum rate at which to serve the VC. Usually, no upper bounding is necessary and the default r_{\max} is the line rate, L . The CAC must ensure that $\sum r_g \leq L$. However, $\sum (r_g + r_e)$ can be greater than L . This is called bandwidth
10 overbooking.

It is now described how the Dynamic Bitmap Generator scheduler 10 can be enhanced to provide overbookable and guaranteed bandwidths. As shown in figures 8 and 9, the BMG generates two
15 bitmaps, G-bitmap (guaranteed rate bitmap) and the O-bitmap (overbookable rate bitmap) simultaneously. The O-bit (G-bit) of a VC (VC group or VC supergroup) is set when the counter 44 (figure 8) loaded with the overbookable rate interval hits zero before (after) the counter 44 loaded with the guaranteed rate interval hits zero as shown in figure 9. The two bitmaps are put into two
20 separate FIFOs, the G-FIFO and the O-FIFO, respectively. The frame number (or the inter-frame number) is included in the G-bitmap and is used to determine how many of the bits of the O-bitmap is selected at each frame.

The rate-limiting is implemented by having a register
25 which stores the next eligible transmission slot. The next eligible transmission slot is equal to the last transmission slot plus $D_{\min} = 1/r_{\max}$. The current slot number is compared with the next eligible transmission slot and if the current slot number is less, then bitmap is inhibited. Otherwise, the bitmap is set and
30 the next eligible transmission slot is loaded with a value equal to current slot number plus D_{\min} .

WHAT IS CLAIMED IS:

1. A scheduler for a server comprising:

a first level generator associated with groups of connections; and

a second level generator associated with connections corresponding to the groups of connections, said first level generator identifying which connections in the second level generator corresponds to a group in the first level generator that are to be considered for service, said second level generator identifies the connections corresponding to the group to receive service from the server, said second level generator in connection with said first level generator.

2. A scheduler as described in Claim 1 including a first level filter mechanism which filters out inactive groups of connections, said first level filter mechanism connected to the first level generator and the second level generator.

3. A scheduler as described in Claim 2 including a second level filter mechanism which filters out inactive connections, said second level filter mechanism connected to the second level generator.

4. A scheduler as described in Claim 3 including a zero level generator associated with supergroups corresponding with groups, said zero level generator in connection with the first level generator, said zero level generator identifying which groups in the first level generator correspond to a supergroup in the zero level generator that are considered for service.

5. A scheduler as described in Claim 4 including a zero level filter mechanism which filters out inactive supergroups, said zero level filter mechanism connected to the zero level generator and the first level generator.

6. A scheduler as described in Claim 5 wherein the zero level generator includes a zero level bitmap generator which generates a zero level schedule bitmap which indicates the supergroup to be scheduled for service, the first level generator includes a first level bitmap generator which indicates the group to be scheduled for service, and the second level generator includes a second level bitmap generator which generates a second level schedule bitmap which indicates the connections to be scheduled for service.

7. A scheduler as described in Claim 6 wherein the zero level, first level and second level filter mechanism includes a zero level filter encoder, first level filter encoder and second level filter encoder, respectively, which filters out inactive supergroups from the zero level schedule bitmap and encodes the zero level schedule bitmap with inactive supergroups removed, which filters out inactive groups from the first level schedule bitmap and encodes the first level schedule bitmap with inactive groups removed, and which filters out inactive connections from the second level schedule bitmap and encodes the second level schedule bitmap with inactive connections removed, respectively.

8. A scheduler as described in Claim 7 including an interface which maintains a zero level active bitmap, a first level active bitmap and a second level active bitmap having only active connections corresponding to the zero level schedule bitmap, first level schedule bitmap and second level schedule bitmap, respectively.

9. A scheduler as described in Claim 8 wherein the zero level filter encoder reads the zero level schedule bitmap and ANDS it with the zero level active bitmap to filter out inactive supergroups, the first level filter encoder reads the first level schedule bitmap and ANDS it with the first level active bitmap to filter out inactive groups, and the second level filter encoder reads the second level schedule bitmap and ANDS it with the second level active bitmap to filter out inactive supergroups.

10. A scheduler as described in Claim 9 wherein the zero level bitmap generator, first level bitmap generator and second level bitmap generator dynamically generates bits for each supergroup, group and connection, respectively.

11. The scheduler as described in Claim 10 wherein the zero level bitmap generator includes a counter for each supergroup which is decremented as a function of an intercell interval, wherein the intercell interval is the time it takes for the server to service a cell, the first level bitmap generator includes a counter for each group which is decremented as a function of the intercell interval, and the second level bitmap generator includes a counter for each connection which is decremented as a function of the intercell interval.

12. A scheduler as described in Claim 11 wherein the zero level bitmap generator sets a bit for a supergroup whose counter decrements to zero, the first level bitmap generator sets a bit for a group whose counter decrements to zero, and the second level bitmap generator sets a bit for a connection whose counter decrements to zero.

13. A scheduler as described in Claim 12 wherein the zero level bitmap generator, first level bitmap generator and

second level bitmap generator each include a rate limiting counter associated with each counter, wherein the bit for the supergroup, group or connection, respectively, is set whenever both the counter and the corresponding rate limiting counter decrements to zero.

14. A scheduler as described in Claim 13 wherein the zero level bitmap generator, first level bitmap generator and second level bitmap generator each generate a guaranteed rate bitmap for supergroups, groups and connections, respectively, which receive service before any other supergroups, groups or connections, respectively, in the respective schedule bitmaps.

15. A scheduler as described in Claim 14 wherein the zero level bitmap generator, first level bitmap generator and second level bitmap generator proportionately reduce the service to each supergroup, group and connection, respectively, when overbooking occurs and when total bandwidth is less than the line rate.

16. A method for scheduling service of a server comprising the steps of:

identifying a group of connections with a first level generator to receive service from the server; and

identifying connections corresponding with the group of connections with a second level generator to receive service from the server.

17. A method as described in Claim 16 including after the identifying the group of connections step, there is the step of filtering out inactive groups of connections in regard to the first level generator.

18. A method as described in Claim 17 including after the identifying the connections step, there is the step of filtering out inactive connections in regard to the second level generator.

19. A method as described in Claim 18 including before the step of identifying the group of connections, there is the step of identifying groups in the first level generator corresponding to a supergroup and a zero level generator.

20. A method as described in Claim 19 including after the identifying groups step, there is the step of filtering out inactive supergroups of connections in regard to the zero level generator.

21. A method as described in Claim 20 wherein the filtering out the inactive supergroups step includes the step of ANDing a zero level schedule bitmap of the zero level bitmap generator with a zero level active bitmap of an interface to filter out inactive supergroups.

22. A method as described in Claim 21 wherein the filtering out the inactive groups step includes the step of ANDing a first level schedule bitmap of the first level bitmap generator with a first level active bitmap of an interface to filter out inactive groups.

23. A method as described in Claim 22 wherein the filtering out the inactive connections step includes the step of ANDing a second level schedule bitmap of the second level bitmap generator with a second level active bitmap of an interface to filter out inactive connections.

24. A method as described in Claim 23 wherein the identifying the groups of connections includes the step of generating dynamically the zero level schedule bitmap, the identifying the group step includes the step of generating dynamically the first level schedule bitmap, and the identifying the connections step includes the step of generating dynamically the second level generator schedule bitmap.

25. A method as described in Claim 24 wherein the step of generating the zero level schedule bitmap includes the step of decrementing a counter for each supergroup every intercell interval; the step of generating the first level schedule bitmap includes the step of decrementing a counter for each group every intercell interval; the step of generating the second level schedule bitmap includes the step of decrementing a counter for each connection every intercell interval.

26. An apparatus for serving connections comprising:

a server;

a memory in which data of the connections is stored, said memory connected to the server; and

a hierarchical scheduler which schedules when the data of the connections in the memory is to receive service from the server, said scheduler connected to said server and said memory.

27. An apparatus for serving connections comprising:

a server;

a memory in which cells of the connections are stored, said memory connected to the server; and

a scheduler which schedules when the cells of the connections in the memory are to receive service from the server based on intercell intervals, wherein an intercell interval is how long the server takes to service a cell, said scheduler connected to said server and said memory.

28. An apparatus as described in Claim 27 wherein the intercell intervals are inversely proportional to bandwidth allocated to a connection.

29. An apparatus as described in Claim 27 wherein spacing at intercell intervals of cells is performed by either statically storing a set of schedule bitmaps or by dynamically generating the schedule bitmap specifying which connections are to be served.

30. An apparatus as described in Claim 12 wherein each counter at each level has a different number of bits.

31. An apparatus as described in Claim 8 wherein each active bitmap has a bit which is set to 1 when an associated connection is active and is set to 0 when an associated connection is inactive.

32. An apparatus as described in Claim 1 wherein connections arise from entities, and including multiple counters associated with each entity which have multiple bits, including multiple schedule bitmaps associated with each entity that are used to schedule connections from the corresponding entity at different priorities or a combination of priorities.

33. An apparatus for serving connections comprising:

a server;

a memory in which cells of the connections are stored,
said memory connected to the server; and

a scheduler having a schedule bitmap which schedules when
cells of the connections in the memory are to receive service from
the server, said scheduler connected to said server.

34. An apparatus for serving connections comprising:

a server;

a memory in which cells of the connections are stored,
said memory connected to the server; and

a scheduler having a schedule bitmap which is either
statically stored or dynamically generated which schedules when
cells of the connections in the memory are to receive service from
the server, said scheduler connected to said server.

35. An apparatus for serving connections comprising:

a server;

a memory in which cells of the connections are stored,
said memory connected to the server; and

a scheduler which maintains active bitmaps which indicate
which connections are active which schedules when cells of the

connections in the memory are to receive service from the server, said scheduler connected to said server.

36. An apparatus for serving connections comprising:

a server;

a memory in which cells of the connections are stored, said memory connected to the server; and

a scheduler having a schedule bitmap and active bitmaps which indicate which connections are active, the scheduler filters out inactive connections from the schedule bitmap by ANDing schedule bitmap with the active bitmaps, the scheduler schedules when cells of the connections in the memory are to receive service from the server, said scheduler connected to said server.

37. An apparatus for serving connections comprising:

a server;

a memory in which cells of the connections are stored, said memory connected to the server; and

a scheduler having schedule bitmaps which can contain multiple bits per connection to schedule different types of bandwidth, the scheduler schedules when cells of the connections in the memory are to receive service from the server, said scheduler connected to said server.

38. An apparatus for serving connections comprising:

a server;

a hierarchical scheduler having levels of hierarchy, the scheduler can enforce rate limiting at each level of the hierarchy, the scheduler schedules when cells of the connections in the memory are to receive service from the server, said scheduler connected to said server.

ABSTRACT OF THE DISCLOSURE

METHOD AND APPARATUS FOR DYNAMIC
BITMAP GENERATOR SCHEDULER

A scheduler for a server. The scheduler includes a first level generator associated with groups of connections. The scheduler includes a second level generator associated with connections corresponding to the groups of connections. The first level generator identifying which connections in the second level generator corresponds to a group in the first level generator that is to be considered for service. The second level generator identifies the connections corresponding to the group to receive service from the server. The second level generator in connection with the first level generator. A method for scheduling service of a server. An apparatus for serving connections. The apparatus includes a hierarchical scheduler connected to the server. An apparatus for serving connections. The apparatus includes a scheduler connected to the server which schedules based on intercell intervals.

[illegible]




FIG 3

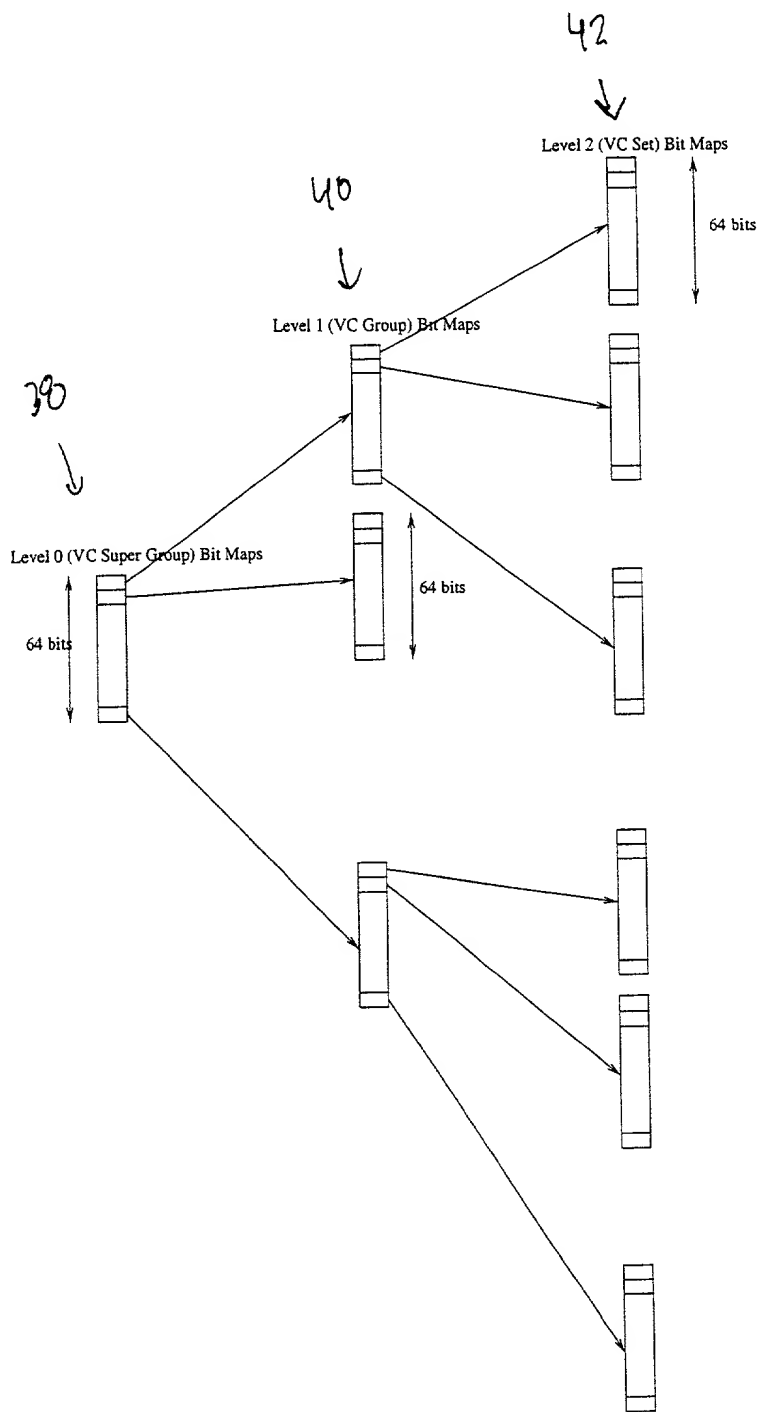


Fig 4




FIG 5

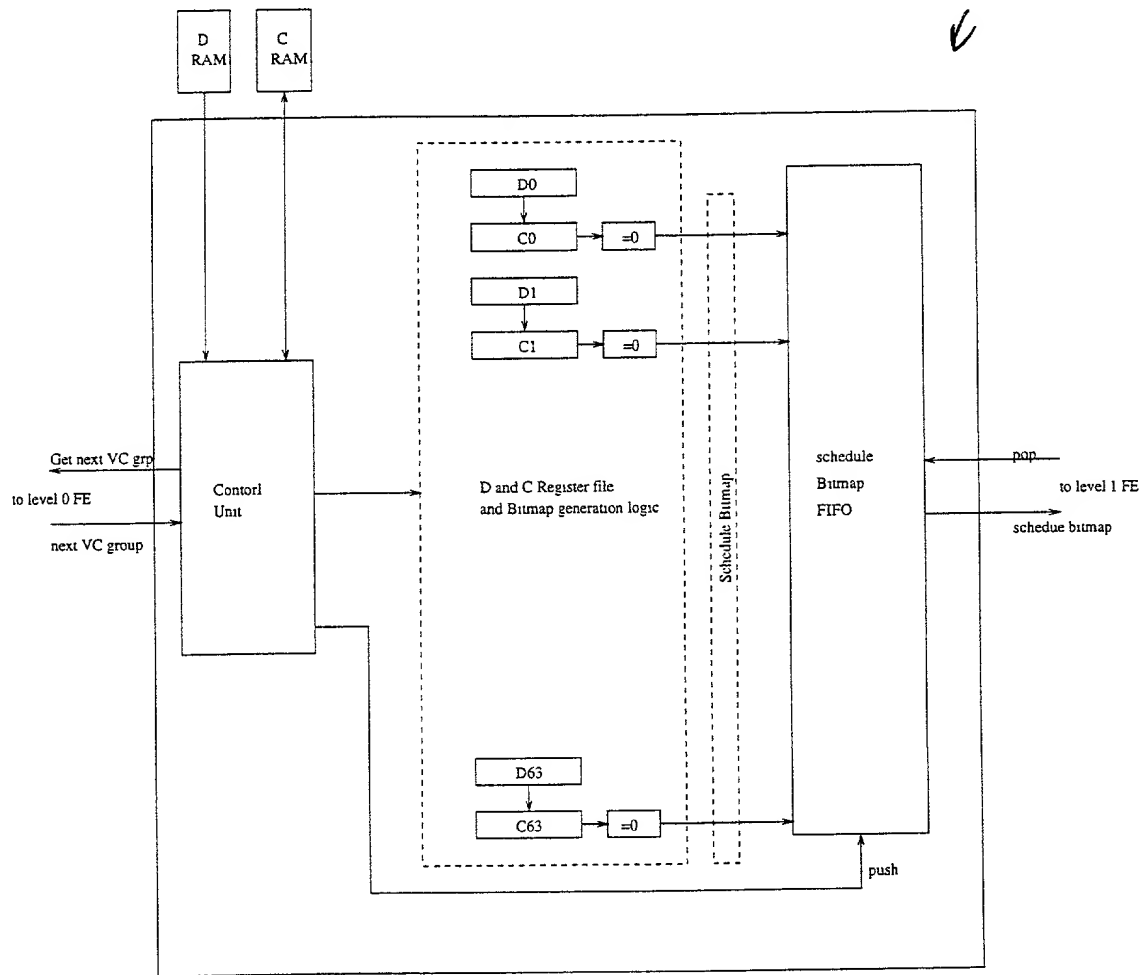


FIG 6




FIG 7




Fig 8




Fig 9




FIG 10

47




FIG 11




FIG 12

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND APPARATUS FOR DYNAMIC BITMAP GENERATOR SCHEDULER

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on _____ as

Application Serial No. 0 /

and was amended on _____
(if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, §1.56(a).

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

Priority Claimed

| | | | | |
|-------------------|--------------------|---------------------------------|---------------------------------|--------------------------------|
| _____ (Number) | _____ (Country) | _____ (Day/Month/Year Filed) | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| _____ (Number) | _____ (Country) | _____ (Day/Month/Year Filed) | <input type="checkbox"/> Yes | <input type="checkbox"/> No |
| _____ (Number) | _____ (Country) | _____ (Day/Month/Year Filed) | <input type="checkbox"/> Yes | <input type="checkbox"/> No |

I hereby claim the benefit under Title 35, United States Code, §120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, §1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

| | | |
|--------------------------|---------------|--|
| 0 / | | |
| (Application Serial No.) | (Filing Date) | (Status) (patented, pending, abandoned) |
| 0 / | | |
| (Application Serial No.) | (Filing Date) | (Status) (patented, pending, abandoned) |

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. (list name and registration number)

Ansel M. Schwartz, Reg. No. 30,587

Send Correspondence to:

Ansel M. Schwartz

412/621-9222

Direct Telephone Calls to: (name and telephone number)

| | |
|---|--------------|
| Full name of sole or first inventor | |
| Meenarachagan Vishnu | |
| Inventor's signature | Date |
| <i>Meenarachagan</i> | Feb 10, 2000 |
| Residence | |
| 117 Lexington Drive, Cranberry, PA 16066 | |
| Citizenship | |
| Canada | |
| Post Office Address | |
| 117 Lexington Drive, Cranberry, PA 16066 | |
| Full name of second joint inventor if any | |
| Second inventor's signature | Date |
| Residence | |
| Citizenship | |
| Post Office Address | |

(Supply similar information and signature for third and subsequent joint inventors.)